

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
25 November 2004 (25.11.2004)

PCT

(10) International Publication Number
WO 2004/102438 A2

(51) International Patent Classification⁷: **G06F 17/60**

(21) International Application Number:
PCT/EP2004/005211

(22) International Filing Date: 14 May 2004 (14.05.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/471,237 16 May 2003 (16.05.2003) US

(71) Applicant (for all designated States except US): SAP
AKTIENGESellschaft [DE/DE]; Neurtstrasse
16, 69190 Walldorf (DE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): SCHMIDT, Patrick
[DE/DE]; Kaiserstrasse 63, 69115 Heidelberg (DE).
BAEUEERLE, Stefan [DE/DE]; Im Bangert 13, 69254
Malsch (DE). NAUNDORF, Stephan [DE/DE]; Brechtel-
strasse 4, 69126 Heidelberg (DE). BURGMEIER,
Hermann [DE/DE]; Konrad-Adenauer-Ring 15, 69126
Nussloch (DE). LIEBIG, Christoph [DE/DE]; Schlatter-
strasse 2, 68542 Heddeshelm (DE).

(74) Agent: ROCKE, Carsten; Müller-Boré & Partner, Grafinger
Strasse 2, 81671 München (DE).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN,
CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI,
GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE,
KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD,
MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG,
PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM,
TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM,
ZW.

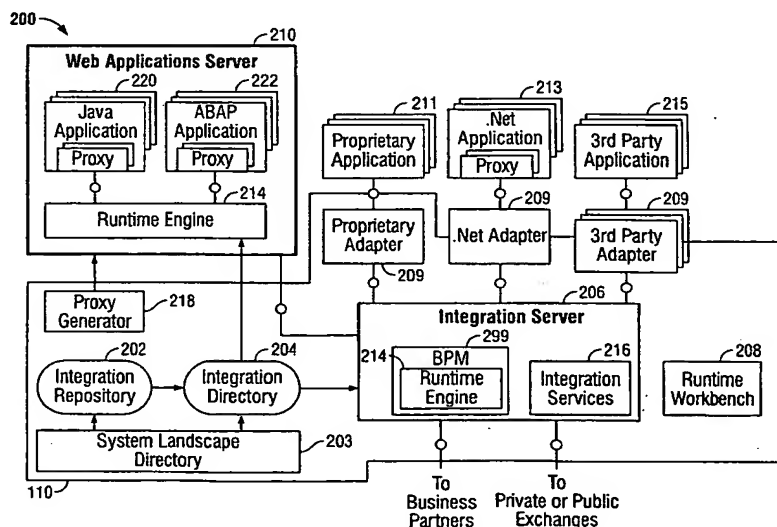
(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,
FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI,
SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished
upon receipt of that report

[Continued on next page]

(54) Title: BUSINESS PROCESS MANAGEMENT FOR A MESSAGE-BASED EXCHANGE INFRASTRUCTURE



(57) Abstract: Methods and systems for managing integration of a heterogeneous application landscape are disclosed. The landscape is defined by one or more business process. A business process management system includes an integration server connected between two or more applications in the landscape. The integration server includes a business process engine configured to execute one or more business processes that define message-based interactions between the two or more applications. The business process management system further applications. The business process management system further includes a runtime engine, under direction of the business process engine, for executing one or more messaging services on the message-based interactions between the two or more applications.



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

**BUSINESS PROCESS MANAGEMENT FOR
A MESSAGE-BASED EXCHANGE INFRASTRUCTURE**

CROSS REFERENCE TO RELATED APPLICATIONS

5

[0001] This application claims priority to U.S. Provisional Patent Application No. 60/471,237, filed on May 16, 2003.

BACKGROUND

10 [0002] The following description relates to enterprise systems and associated architectures and techniques for collaborative business processes. Companies face an increasing need for integration of and collaboration among their information and enterprise software systems. In most current system landscapes, many components are directly
15 connected in a one-to-one relationship with other components, with the integration capabilities hardwired into the application components and individual mappings programs. Under these conditions, managing the collaborative sharing of information is difficult. These infrastructures can
20 rarely represent actual business processes accurately, and are limited in their flexibility to dynamic business scenarios that govern business processes.

[0003] New electronic business collaboration, however, typically requires connectivity among all applications
25 inside and outside of company boundaries. Networks such as the Internet provide opportunities for systems to communicate almost instantly with other systems or individuals. Business processes that once were restricted to intranets and their users are now moving to the Internet
30 to become an effective composition of Web services. A Web service is a programmable, self-contained, self-describing,

modular application function that can be published,
discovered or invoked through an open Internet standard.
However, comprehensive system upgrades of existing
enterprise software, or large-scale replacement strategies
5 in heterogeneous system landscapes tend to be too costly or
otherwise and simply unfeasible in terms of time and capital
resource costs.

[0004] While technical connectivity is provided using
open protocols and standards like the hypertext transfer
10 protocol (HTTP) and extensible markup language (XML), the
challenge of mapping different business semantics remains.
To capture future rounds of efficiency gains, enterprises
increasingly will be required to deploy a new breed of
collaborative business processes that cross enterprises or
15 functions within an enterprise. In addition, enterprises
will increasingly need to process and manage real-time
scenarios instead of performing batch processing.

SUMMARY

[0005] This document discloses a business process
20 management (BPM) system and method for managing and
executing message-based communication according to actual
business processes and scenarios. The BPM system includes a
business process engine (hereinafter also referred to as
"process engine" for simplicity) configured to execute
25 instructions according to business processes of an
application system landscape. The process engine keeps a
state on the integration server to handle correlated
messages. Correlation means that the process engine is able
to identify messages as being semantically linked, such as a
30 purchase order and an invoice regarding a single product.

The process engine is able to receive and hold multiple messages in one process instance.

[0006] The process engine is configured to execute an explicit serialization (including handling of
5 acknowledgements). This serialization is independent from sender and receiver system(s), and works when more than one receiver and sender system is involved, as distinct from an exactly once in orderEOIO process. The process engine is further configured to execute a process-controlled
10 multicast, to include handling of response messages.

[0007] The process engine is also able to collect correlated messages. The process engine also provides transformation services, which are configured for merging multiple messages into one message (such as a purchase order
15 header and an arbitrary number of purchase order positions into a single purchase order), as well to split a message into several parts, e.g. splitting a purchase order into a purchase order header and into an arbitrary number of purchase order positions. A business process (also referred
20 to hereinafter as simply a "process") may use information from all participating messages (messages that the process received earlier), and can route responses to the original sender of an earlier request.

BRIEF DESCRIPTION OF THE DRAWINGS

25 [0008] These and other aspects will now be described in detail with reference to the following drawings.

[0009] FIG. 1 is a simplified block diagram of an exchange system for integrated, message-based collaboration.

[0010] FIG. 2 is a block diagram of an exchange infrastructure.

[0011] FIG. 3 is a detailed block diagram of an integration repository, integration directory, and runtime engine for collaborative processing.

[0012] FIG. 4 is a block diagram illustrating a process for communicating a single message between two or more applications.

[0013] FIG. 5 is an architectural block diagram of a BPM system including an integration server and a business process engine.

[0014] FIG. 6 is a workflow diagram of a BPM system.

[0015] FIG. 7 illustrates links to and from business processes.

[0016] FIG. 8 shows a serialization use case of a BPM system.

[0017] FIG. 9 shows a basic merge scenario of a BPM system.

[0018] FIG. 10 illustrates a stateless communication scenario in a BPM system.

[0019] FIG. 11 illustrates a stateful communication scenario in a BPM system.

[0020] FIG. 12 shows three types of communication scenarios in a BPM system.

[0021] FIG. 13 illustrates the use of context objects to distinguish from among different send steps on the same interface.

5 [0022] FIG. 14 illustrates an aspect of a merge case in a BPM system.

[0023] FIG. 15 illustrates a value mapping process as a process-normalized value-format operation.

[0024] FIG. 16 illustrates an embodiment of a process engine runtime.

10 [0025] FIG. 17 illustrates an implicit cancellation operation.

DETAILED DESCRIPTION

[0026] The systems and techniques described here relate to management of business processes that define a message
15 communication protocol between applications in a heterogeneous system landscape. The business process management system and method is optimally implemented in an exchange infrastructure configured to integrate and drive collaboration between various applications in the landscape
20 using open standards and transport protocols such as XML and HTTP.

[0027] FIG. 1 is a simplified block diagram of a system
100 for integration and message-based interaction of applications. The system 100 includes an exchange
25 infrastructure (XI) 110 for collaborative processing among internal components (ICs) 102 of an enterprise, and between external components (ECs) 104 that communicate to one or

more ICs 102 through a firewall 105. The ICs and ECs 102 and 104 represent any of a number of processes or services and their software and hardware, such as Web portals, buying or selling programs, electronic mail, business management
5 programs, project planning programs, etc., and are preferably Web-based applications. Each of the ICs/ECs 102, 104 communicates via messaging with one or more other components according to at least one of a number of communication protocols or standards.

10 [0028] The XI 110 is a self-contained, modularized exchange platform for driving collaboration among the components 102, 104. The XI 110 includes a central integration repository and directory storing shared collaboration knowledge. The XI 110 supports open standards
15 such as various standard markup languages like the extensible markup language (XML), web service description language (WSDL), and simple object access protocol (SOAP) to provide an abstraction of technical interfaces for the components 102, 104, and for message-based communications
20 across heterogeneous component interfaces. The self-contained, modularized functions of the XI 110 can be provided as one or more Web services based on standard Internet technology, and therefore can be published, discovered, and accessed within a network of components 102,
25 104 using open standards.

[0029] FIG. 2 illustrates a system landscape 200 including an XI 110 for facilitating message-based collaboration among applications. The exchange
infrastructure 110 includes an integration repository 202,
30 an integration directory 204, a system landscape directory

203, and an integration server 206. The integration repository 202 captures design-time collaboration descriptions of all software components that can communicate via the XI 110. The integration directory 204 captures configuration-specific collaboration descriptions of the system landscape 200 at runtime, which includes accessing actual component installations from the system landscape directory 203 and connectivity descriptions for external components, all of which represents the shared business semantics of the system landscape 200. The integration server 206 uses the shared business semantics at runtime to execute message-based collaboration among the active software components.

[0030] The integration server 206 includes a runtime engine 214 that provides messaging and business process control at runtime for connecting services and managing the process flow of value chains. The BPM system 299 resides on the integration server 206 below the runtime engine 214. The BPM system 299 includes a business process engine, or "process engine," (not shown) that runs on top of the runtime engine 214 for runtime execution of business process management.

[0031] The integration server 206 also includes integration services 216 that typically require an application-specific implementation. Like the integration repository 202 and integration directory 204, the integration server 206 is configured for deployment within any existing system infrastructure. The integration server 206 is preferably a dedicated server that applies the shared collaboration knowledge of the integration directory 204 of

the supported system landscape in a runtime collaboration environment. A runtime workbench 208 allows organizations or users to manage the reliable operation of the XI 110.

[0032] The XI 110 also includes various adapters 209 that
5 provide connectivity between the integration server 206 and
proprietary applications 211, Web-based services 213, and
third party applications 215. The XI 110 can also include
Web applications server 210 that provides Web-based
applications programmed according to standard computing
10 platforms using web-specific programming languages such as
Java and ABAP, for instance. The Web applications server
210 also includes an instance of the runtime engine 214 for
providing messaging and business process control between
Web-based applications such as Java applications 220 and
15 ABAP applications 222, and other components.

[0033] New interfaces for software components can be
defined using an application component employing a proxy,
which allows the interface for the software component to be
implemented locally in the XI 110. Proxies make the
20 communication technology stack transparent to applications,
and present an application with a programming language-
dependent interface. The proxies can be generated by a
proxy generator 218 based on information stored on the
integration repository 202. The proxy generator 218 uses
25 the interface information described via a standard Web-based
language such as WSDL and XSDL to create platform- and
programming language-dependent code in the application
development system. The communication logic can be
implemented based on the proxy that represents the interface
30 description of the respective development platform, such as

Java, ABAP, and .NET for the web-based applications 213. The proxies convert platform-specific data types into XML and provide access to the component-specific local integration engine. On the outbound side, proxies are
5 generated completely. Outbound proxies can be called via a service invocation provided by an application's developer. On the inbound side, only proxy skeletons need to be generated, as implemented by the receiving application.

[0034] FIG. 3 illustrates the integration repository 202,
10 the system landscape directory 203, the integration directory 204 and an instantiation of the runtime engine 214 in greater detail. The integration repository 202 includes design-time business processes 232, routing objects 234, mappings 236, and interfaces 238, all of which are defined
15 according to one or more business scenarios 230. The integration repository 202 accesses descriptions of all software components 240 in the system landscape from the system landscape directory 203. The business scenarios 230 of the integration repository 202 describe and configure
20 message-based interaction between application components or enterprises. An enterprise can select one or more business scenarios described in the integration repository 202 as a best practice for rapid configuration of the XI 110.

[0035] The business processes 232 can be implemented as
25 extensible compound Web services executed using a business process engine 274. Each business process 232 is modeled centrally in the integration repository 202. A company or user designs each business process 232 according to its business needs, independently of the technical
30 implementation. There may be several categories of business

process templates: i.e. generic business processes, industry-specific processes, and company-specific processes, for example. Each process identifies the Web services that are needed and that must be interconnected. In one specific
5 implementation, business processes 232 are defined using a graphical interface. The graphical business process definition language is defined according to the BPEL4WS 1.1 specification. Business processes 232 can be exported into and imported from a standardized format (this is today the
10 BPEL4WS 1.1 specification). An extensible import/export framework makes it possible to provide export/export facilities for other standards or new versions of BPEL4WS. The business process engine 274 can then interpret these models and execute them to drive collaboration among
15 software components.

[0036] Routing objects 234 are predefined criteria to determine potential receivers of messages that must be distributed between components and business partners during collaborative processing. Information about the routing
20 objects is used for receiver determination to avoid having to process a complete message before distribution. Mappings 236 define required transformations between message interfaces 238, message types, or data types in the integration repository 202. These transformations cover
25 structural conversions and value mappings. Structural conversions are used for semantically equivalent types that are syntactically or structurally different, whereas value mapping may be used when an object is identified by different keys in multiple systems. In a specific
30 implementation, a graphical mapping tool is provided to assist in mapping, and transforming data is based on the

Extensible Stylesheet Language Transformation (XSLT) or Java code.

[0037] The integration repository 202 is the central point of entry for interface development, storage and retrieval, and includes interfaces 238 that describe all message interfaces of all software components in the system landscape. Accordingly, the interfaces 238 can be implemented on any software component using any technology. Message interfaces are made up of message types, which are in turn made up of data types. The data types can be described using XML Schema Definition Language (XSDL). An example of a data type is "address," which is used in the message type "Create PO" and can be reused for the message type "Create Invoice." Interfaces 238 can be arranged according to any classification, such as inbound, outbound and abstract, or synchronous and asynchronous.

[0038] The components 240 represent component descriptions that include information about application components, as well as information relating to their dependencies on each other. In a specific implementation, the component descriptions are based on the standard Common Information Model (CIM) of the Distributed Management Taskforce. Since the integration repository 202 includes design-time information, only component-type information, independent of actual installation, is stored as components 240 in the system landscape directory 203. The component descriptions can be added using an API or interactively using a graphical user interface.

[0039] The integration directory 204 details information from the integration repository 202 that is specific to the

configuration of each component as installed in the system. The configuration-specific collaboration descriptions of the integration directory 204 can be generated automatically from content in the integration repository 202 or manually
5 by a user using a graphical user interface. In one implementation, the integration directory 204 is built on a Java platform and its content is represented via XML using open Internet standards. The integration repository 202 can be upgraded without affecting the integration directory 204
10 or any runtime collaborative processes. The user then decides which changes should be transferred to the integration directory 204, either as predetermined automatic upgrades or manually via graphical tools.

[0040] The integration directory 204 includes
15 configuration-specific descriptions of business scenarios 250, business processes 252, context objects 254, and executable mappings 256. The integration directory 204 also includes descriptions of active Web services 258, and active business partners 260. The integration directory 204 uses a
20 description of the active system landscape 262 from the system landscape directory 203. The business scenarios 250 in the integration directory 204 represent the overall view of the interaction among interfaces and mappings 256 in the context of the actual configuration relevant for the
25 specific implementation. The business processes 252 represents an executable description of all active business processes.

[0041] The context objects 254 determine the receivers of a message on a business level. In one specific
30 implementation, the content of a message is used as a

context object 254. Other parameters may also be used. Relevant input parameters include the sender, the sender message type, the message to identify the receivers, and the receiver message type. The context object 254 can be
5 described declaratively using XML Path Language (Xpath, i.e. by using a graphical tool) or can be coded in Java. The integration engine 214 at runtime accesses information on the context object 254.

[0042] The context objects 254 may use logical terms to
10 describe senders and receivers in order to separate them from the physical address provided by the Web services 258 described in the integration directory 204. The physical address can therefore be changed without changing business-oriented content. Mappings 256 in the integration directory
15 204 represent mappings required in the active system landscape, in contrast to the integration repository mappings 236 that contains all supported mappings. Some new entries however, such as a new sequence of mappings, can be made only in the integration directory 204 to address
20 additional Web services for mapping, for example. The integration engine 214 accesses the integration directory mappings 256 at runtime.

[0043] Context objects 254 provide a unique name for
accessing semantically identical payload information. For
25 instance, a context object can provide a unique access name for 'plant' for invoice and purchase order. The XPath for 'plant' in an invoice can be defined as '/A/B/C/plant' and the XPath for 'plant' in a purchase order looks like
'X/Y/Z/werk'. The context object 254 'plant' is assigned to
30 the message interface invoice and purchase order where the

XPaths as above mentioned are specified. This makes sure that the XPath for plant is not defined at n different places.

[0044] Web services 258 describe interfaces implemented within the current active system landscape, as well as active Web services supported by described business partners 260. As such, information describing Web services 258 can be exchanged with Universal Description, Discovery, and Integration (UDDI) compatible directories or added manually. Each Web service 258 description also provides physical addressing details, access information, and other special attributes such as uniform resource locator (URL), protocol, and security information. In one implementation, the Web services 258 are described in WSDL, and SOAP and ebXML are used as messaging protocols. The integration engine 214 accesses information about the Web services 258 at runtime as well.

[0045] The system landscape 262 of the system landscape directory 203 describes the current system landscape that uses the XI 110. The system landscape 262 describes the components that are installed and available on certain machines within the system, the instance or client that was chosen, further information on the installed components, other system landscapes, and so on. The system landscape 262 description is based on an open architecture and can adhere to any widely accepted standard such as the Common Information Model (CIM). Thus, many proprietary and third party components can be configured to automatically register themselves in the system landscape 262 upon being installed within the actual system landscape. Access interfaces to

the system landscape 262 description can be based on open standards as well, such as the Web-based Enterprise Management (WBEM) and SOAP standards.

[0046] Business partners 262 defines information for
5 business partners of an enterprise, such as names, addresses, and URLs, but may also contain more detailed and sophisticated information. For instance, the business partners 262 may include a description of the message formats that can be directly received and processed, or of
10 security protocols used for safe communications, or trading terms that are employed in the partnership. The kind of information stored in business partners 262 can be governed by enterprise-specific decisions of the enterprise using the XI 110.

15 [0047] The integration directory 204 and the runtime engine 214 form a collaborative runtime environment for executing collaborative business processes. The collaborative runtime environment provides all runtime components relevant for exchanging messages among the
20 connected software components and business partners. The integration server 206 executes the collaborative runtime environment or Web application server 210, either of which can include an instance of the runtime engine 214 in accordance with informational resources provided by the
25 integration directory 204.

[0048] The runtime engine 214, which exchanges all messages between the various interconnected components, includes two layers: an integration layer 272 and a messaging and transport layer (MTL) 280. The integration
30 layer 272 includes a business process engine 274 executing

centrally modeled business processes, a logical routing service 276 and a mapping service 278. The MTL 280 provides a physical address resolution service 282, a messaging and queuing service 284, a transport service 286 via HTTP, and a database 288. The integration services 216 in the integration server 206 can support the runtime engine 214. An MTL 280 is also included in each instantiation of the runtime engine 214 in Web applications servers 210, as well as in each adapter 209 of the adapter framework connecting to various software components. Each MTL 280 has a role in the execution of the EO protocol, as will be explained further below.

[0049] At runtime, business processes 252 are instantiated and executed by the business process engine 274, which executes the respective Web services described in Web services 258 independent of their location according to the business process model. The business process engine 274 is independent of the semantics of the executed business processes 252, and is configured as a mediator and facilitator for business processes 252 to interact with technical components of the runtime system landscape.

[0050] FIG. 4 is a block diagram illustrating several functions of the runtime engine 214 and BPM 299 in a process of exchanging a message between applications. A sending application 303 resides in a sending component system 302, which represents the hardware and software platform of the sending application 303. One or more receiving applications 305 each reside in a receiving component system 304. A communication path for a message 310 can include an outbound proxy 307 at the outbound interface from the sending

component system 302, through the runtime engine 214 and adapter 309 to the receiving component system 304. A receiving component system 304 may also utilize an inbound proxy 311 rather than an adapter. The configuration and connectivity of the shown receiving component systems 304 is merely exemplary, and it should be noted that such configuration and connectivity could take any number of forms. The pictured example illustrates both asynchronous and synchronous communication. In synchronous communication, routing and physical address resolution is only needed for the request as the response is transferred to the sender, which is already known.

[0051] For a given message the logical routing service 276 uses information on the sending application and the message interface to determine receivers and required interfaces by evaluating the corresponding routing rules, as shown at 312. The routing rules are part of the configuration-specific descriptions of the runtime system landscape provided by the integration directory 204, and can be implemented as XPath expressions or Java code. The mapping service 278 determines the required transformations that depend on message, sender, and sender interface, as well as the receiver and receiver interface, at 314. In the case of asynchronous communication, even the message direction is determined to appropriately transform input, output, and fault messages.

[0052] After retrieving the required mapping from the integration directory 204, the mapping service 278 can either execute XSLT mappings or Java code (or any combination in a given sequence) to the content of the sent

message. Below the integration layer, messaging, queuing, and transport services 284 move the message to the intended or required receiver(s). After the message is transformed into the format expected by each receiver, the physical
5 address of the required receiver service and other relevant attributes are retrieved from the integration directory 204 and mapped to the message, at 316.

[0053] A queuing engine (not shown) in the messaging and queuing service 284 stores ingoing, outgoing, erroneous, and
10 work-in-progress messages persistently. The messaging layer of the runtime engine 214 provides queuing functions for the physical decoupling of application components and guarantees messages are delivered exactly once according to a protocol (i.e. the "EO protocol"). The transport service 286 enables
15 the runtime engine 214 to act as both a client and server. The transport service 286 implements a client that enables outbound communication and a server that handles inbound communication by accepting incoming documents. Additional server functions can address situations in which the
20 receiver has no server by supporting polling over the transport protocol used. Preferably HTTP is used, but other transport protocols may be used as well.

[0054] FIG. 5 is an architectural block diagram of a BPM system 500, which includes a process engine 504 integrated
25 in an integration server 502. The process engine 504 and integration server 502, as they are called in their runtime configurations, are also respectively known as a process editor and an integration builder in their "definition time" configurations. Process definition 506 and BPM runtime 508
30 in the BPM system 500 are based on different development

platforms. For instance, the process definition 506 is based on Java, such as a J2EE platform 505, and the runtime 508 is based on ABAP. The BPM system 500 includes monitoring and administration tools 524 on the integration
5 server 502.

[0055] The process definition 506 module utilizes XML objects and correlations to define processes, based on deployment rules imported from XI objects 512 from the integration directory 514. The XI objects 512 are based on
10 the routings and mappings defined for the system runtime configuration 516. The XI objects 512 are also used to define business processes 518 in the integration repository 522, and the design-time configuration 520 of the system landscape. Business processes 518 are integrated with all
15 other integration repository 522 objects and tools, which allows links to and from other XI objects. Processes (i.e. patterns and templates) can be delivered to customers, and extension concepts may be provided. Application specific content can also be delivered. Some applications can create
20 processes and/or extensions that can be delivered to customers. The BPM system 500 includes an import/export framework 526 that imports and exports standards-based adapters for universal connectivity. The BPM system 500 can include an interface for receiving user-specified business
25 process details.

[0056] Business process modeling scenarios are also known as modeling patterns (or simply, "patterns"). The following described patterns are high-level building blocks, and can be combined with each other and with atomic process engine
30 504 functions such as deadlines, exceptions, etc.

[0057] 1) Send and Receive: Sending messages controlled by the process engine 504 is often combined with receive steps that wait for a correlated response message. A receive step should wait for the messages starting with the activation of the associated correlation as a queuing mechanism.

[0058] 2) Serialization: This pattern can include the following steps: 1. Receive messages and store them locally in the process data context; 2. Keep the data context and start sending received messages when a certain condition has been fulfilled; and 3. Send received messages in a given order respecting dependencies of receivers. This third step can be: a. Without caring about responses/acknowledgements ("fire and forget"); or b. Receiving a response or an acknowledgement (enables serialization). The process engine 504 can be configured to wait for a technical ACK of or business response from a previously-sent message before sending a next message.

[0059] 3) Transformations/Merge/Split: The process engine 504 transforms messages within the process context. The following transformations can be performed: 1. (N:1) Transform several collected messages to one new message (e.g. transform several invoices to one combined invoice or transform PO header and several PO positions into one PO); 2. (1:N) Transform one message into several other messages (e.g. transform a combined in-voice to invoice respecting the original POs); and 3. (1:1) is a special case of the transformations described above. N:M mappings are also possible if needed.

[0060] 4) Multicast: The process engine 504 can be configured to calculate the receivers of a message (also using content-based conditions) and to send the message to these receivers, either without regard to

5 responses/acknowledgements ("fire and forget") or based on receiving a number of responses/acknowledgements. Messages may be sent out in parallel or sequentially.

[0061] 5) Collect: This pattern uses receive steps in which an arbitrary number of messages can be received. From
10 a process point of view, the end of the collecting scenario can be defined via "push," (i.e. a certain condition is reached, such as N messages have arrived, a certain deadline has been reached, etc.), or "poll" in which the process engine waits for a special message that indicates the end of
15 collecting.

[0062] With reference also to FIG. 5, FIG. 6 illustrates a workflow 600 of a BPM system 500 runtime and respective process engine of the integration server 502 orchestrating several 'client' application systems 503. The integration
20 server 502 is a standalone component that communicates via messages with the client application systems 503. Message-related functions (send, create, transformation, merge, split, etc.) are preferably realized by service calls to messaging layer of the integration server 502 ('lower-level
25 XI-runtime functions'). The process engine 504 preferably does not change the message-payload directly. Rather, messages are changed by transformation, which is explained further below.

[0063] The process engine 504 focuses only processes on
30 the integration server 502, not on application systems 503.

The process engine 504 is not used to control processes in the backend systems, but while it is able to communicate with backend processes via messages, the process engine 504 does not interact with the applications, organizational and user management functions in the backend system(s). The process engine 504 uses the messaging layer a application while Business Workflow uses the application, user- and organizational management of the respective application system. The process engine 504 supports the communication via synchronous outbound interfaces.

[0064] BUSINESS PROCESS MANAGEMENT

[0065] The integration of process management into the XI landscape can described in two parts: 1) the "outside view," which covers all references to and from the business processes; and 2) the "inside view," which looks into the processes and the modeling tools.

[0066] Outside View

[0067] Processes will have representations both in the integration repository 522 and the integration directory 514. Process definitions will be created and stored only in the integration repository 522. This allows the transport of process definitions to the client systems 503. Processes stored in the integration directory 514 will rely on a thin representation, which will point to an associated process definition in the integration repository 522. Each business process, as an XI object (visible in the navigation tree and usable in links from and to other XI objects) will provide the ability to integrate the process engine 502 in the XI environment. Business processes 518 can use

established XI object types, and will not create redundant object types. Business processes 518 include public parts, such as previously-used interfaces, and private parts, which include the process graph using step types and correlations.

5 Process instances can be stopped and restarted in runtime 508. Process instances can also be restarted from a specific step (e.g. if an error occurs during a certain step, restart from that step).

[0068] FIG. 7 illustrates links to and from business
10 processes 518 in the integration repository 522. The links include references to: (2) abstract interfaces 702; (3) context objects 704; and (4) interface mappings 706. Absolute links include: (1) the action of a business scenario references a process definition; (5) an interface
15 mapping 706 references a message mapping 712. Business processes 518 can be used in business scenarios 720, and will act as brokers between business systems.

[0069] A business process 518 owns a process interface 708 which reflects all inbound and outbound communication.
20 Interfaces used in the process interface 708 include two types: process-specific interfaces (a special normalizing interface for the process); and mirrored outbound/inbound interfaces of sending/receiving business systems (to avoid the creation of unnecessary interfaces). Mirroring must be
25 done creating a new abstract interface 702 pointing to the same message type as the original interface. Abstract interfaces 702 can be used in an inbound as well as in an outbound role. That ensures that no mapping for a message received by the process and sent out afterwards without
30 changes (without transformation) is necessary. Would the

process need inbound and outbound messages, a transformation from inbound to outbound would be required. In addition, process-specific interfaces do not need to have proxies in the attached business systems. This leads to the
5 so-called abstract interfaces 702, which are the only type of interfaces that can be used by the business processes 518. Local interfaces may reference other interfaces 708 (to handle the mirroring) and they also may reference message types 710 (to realize process-specific interfaces).

10 **[0070]** Context objects 704 can be used to access payload information via name or other message content. Data may not be written to context objects 704. Interface mappings 706 are addressed by a business process 518 within the transformation step.

15 **[0071]** A business scenario 720 may reference one or more business processes 518. One business process occupies one "swim lane" or process flow. Each process is treated as a business system. Actions within process swim lanes are not stored as separate actions that are reusable. An action
20 represents an interface used by a business process 518 as outbound or inbound interface (or both). In a 'normal' scenario case, not all interfaces of an action must be a target or source of a connection. In the business process case, each action represents one interface with inbound
25 and/or outbound semantics and must be used as target and/or source in a connection. Navigation from the business scenario 720 to the business process 518 is provided.

[0072] FIG. 8 shows a simple serialization use case. O* represents an outbound interface. I* represents an inbound
30 interface. L* represents a local interface. Messages of

interfaces L2 and L4 are received by the business process (each of them can arrive first and start the process) and sent out in sequential order. Each action represents one interface that is used as an inbound interface (receive) and
5 an outbound interface (send) by the business process. Therefore, each action has two connections: to and from the action.

[0073] FIG. 9 shows a basic merge scenario. Messages of interfaces L2 and L4 are received by the business process
10 (each of them can start the process) and transformed into a new message of interface L5 which is sent out. The actions representing the interfaces L2 and L4 have only inbound (receive) connections, while the action representing interface L5 has only an outbound connection (send).

15 [0074] Regarding a business scenario, a business process may be modeled in different ways. One way is called "process first," or bottom up, which includes creating a process definition and process interface. This process signature includes interfaces that have an inbound and/or an
20 outbound role. Messages of these interfaces can be received/sent by a process. The process signature can be derived from the process definition and the container elements and does not need to be defined by the process designer. The inbound part of the process interface is
25 based on receive steps that receive messages of a given interface (i.e. the type of the container element that should be received). The outbound part of the process signature is based on send steps that send messages of a given interface (i.e. type of the container element that is
30 sent)). This method also includes introducing a process

definition in a business scenario. Within the business scenario, actions refer to interfaces of the business process.

- [0075] Another way a process can be modeled is called "scenario first," or top down, which includes the following steps: 1. introduce process definition as a swim lane in a business scenario (create process definition implicitly if it does not exist); and 2. add Actions to the swim lane of the process (implicitly create/edit the process interface).
- [0076] As illustrated in FIG. 7, the transition from integration repository 522 to integration directory 514 should be as effortless as possible. In one embodiment, the process 732 in the integration directory 514 is created, named and linked to the active process definition 518 in the integration repository 522. It is also possible to have more than one process in the integration directory 514 pointing to the same process in the integration repository 522. Thereafter, all other integration directory objects can be linked with the process 732. To ease this transition, processes 732 can be included in a user-configurable transition "wizard" that is displayed in a user interface.

- [0077] The process 732 in the integration directory 514 comprises a thin representation: the process is identified by a name and embraced in a business scenario 730 due to the dependency of routing relations 734 on the business scenario 730. The directory process 732 contains a reference (link) to its originating repository process 518 (active version). Accordingly, the process representation in the integration directory 514 will not contain a process definition.

Accordingly, process customization or configuration in the integration directory 514 is not necessary.

[0078] The process 732 in the integration directory 514 allows specifying routing relations 734 and mapping
5 relations 736 from and to the process 732. That means, mapping and routing functions need to recognize a business process 732 as a source or target (i.e. business processes can be addressed like business systems). It must be visible
10 in an integration builder interface whether a process or a business system is addressed. Processes 732 are also used as senders of messages in the routing. For example, two different send steps should be distinguished within the same process that sends messages of the same interface to different receivers.

15 [0079] The transition from a stateless scenario (without process) to a stateful scenario (with process) should be as smooth as possible. FIG. 10 shows a stateless communication. The connections (routings) O2 to I2 and O4 to I4 can be serialized, where O and I are used interfaces.

20 [0080] FIG. 11 illustrates a stateful scenario. To create a stateful scenario, the connections (routings) O2 to I2 and O4 to I4 need to be removed. The local Interfaces L2 (mirrored interface to O2) and L4 (mirrored interface to O4) as well as the process definition need to be created. Next,
25 a process swim lane (BP1) needs to be introduced in the business scenario. The actions can be taken over from the process interfaces and the connections (routings) from and to the process must be adapted: O2 to L2 and L2 to I2; and O4 to L4 and L4 to I4. Mappings that have been used between
30 interfaces O2 and I2 can be reused between the interfaces L2

and I2 (since L2 is a mirror of O2). The same can be true for O4/L4 and I4.

[0081] The system can include a graphical process modeling tool for graphical modeling of business processes.

5 The graphical process modeling tool includes a process builder graphical user interface (GUI) that graphically illustrates modeling elements and basic modeling paradigms.

[0082] Each process is executed using data that is available within the process. This data is stored in

10 container elements (variables) that are entities of the process container. The container elements can be defined when modeling a process and include a unique name and a type. A container element can be defined as multi-line (i.e. table of elements). Container elements can be grouped
15 according to a type. One such type is interfaces (i.e. for the messages that are handled by the process). Another such type is simple XSD types (i.e. for process control data e.g. counters, etc.) Receivers can also be specified as one data type as needed for multicast.

20 **[0083]** The process interface represents the (interface-) signature of the process. It includes interfaces that have an inbound and/or an outbound role. Messages of these interfaces can be received/sent by a process. The process interface can be derived from the process definition and the
25 container elements. The inbound part of the process interface is based on receive steps that receive messages of a given interface (i.e. type of the container element that should be received). The outbound part of the process interface is based on send steps that send messages of a

given interface (i.e. type of the container element that is sent).

[0084] Messages that are sent to a process need to be delivered not only to the process definition, but also to the correct instance of the process. This "instance routing" can be accomplished generically. Starting a new process instance is a special case. The messages that are handled by a process are loosely coupled. The dependency of messages that should be handled by the same process relies on the business data.

[0085] Correlations are the means to define these dependencies. The declarative specification of a correlation relies on the declarative properties of messages. A property is simply a field within a message. Correlations can be defined using multiple fields. The position of correlation tokens (i.e. values for the correlation properties) in each message can be expressed declaratively in the process definition. Correlations can be instantiated within the scope of a process instance as long as the first message that keeps the information is available. This instantiation is modeled by a send and receive steps.

[0086] Processes can only be triggered by message arrival. Starting a process via application call or via API or via human interaction can be mapped into using triggering messages. Triggering messages can be identified in the receive step. Accordingly, a received message can be declared as a triggering message, similar to BPEL4WS. To be able to declare a receive step as triggering, this step needs to be defined at the beginning of the process.

[0087] Receive steps are used to consume messages that are sent to a process. Each receive step will get its own message, even if multiple receive steps wait for messages of the same interface. If a message is sent to a process and
5 the receive step has not been instantiated yet, the process engine queues the message until the receive step is reached.

[0088] As described above, messages are sent to a process. If no receive step consumes a message, the process
10 will cache the arriving messages. If a receive step is reached, the "oldest" message is fetched out of the cache and the receive step completes. If the process reaches a receive step and the cache is empty, the process waits until a new message arrives. Each receive step instance in the
15 BPM system gets its own message, unlike a classical event-based publish and subscribe, where several event consumers are notified upon the same event.

[0089] The receive step references all correlation(s) which have to be fulfilled to complete the receive step. In
20 the case of more than one correlation, all referenced correlations must be fulfilled to complete the receive step (and-semantics). A receive step also references the correlation(s) that need to be instantiated upon the consumption of the received message. The arrival of a
25 message is the only way in which new data can be brought into the process.

[0090] There can be multiple receive steps for messages of the same interface. As shown in FIG. 12, at least three scenarios may be possible: in sequence; in parallel; and in
30 a loop. These will be discussed now in turn.

[0091] In the sequential scenario, multiple receive steps are waiting for messages of the same interface. The messages arrive at different points in time (t1 and t2) while the receive steps begin a wait period from the correlation activation time. The semantics are as follows: each receive gets its own (different) message; the first message completes the first receive step while the second message completes the second receive step; the first message does not complete both (or more) receive steps. This action avoids race conditions and ensures that no message is lost.

[0092] In the looped scenario, a receive step is waiting within a loop for messages of the same interface. The messages (of the same interface) arrive at different points in time (t1, t2 and t3) while the receive step waits from the correlation activation time on. The semantics are as follows: if no receive step consumes the message, the process will cache all arriving messages; if a receive step is reached, the "oldest" message is fetched out of the cache and the receive step completes; if the process reaches a receive step and the cache is empty, the process waits until a new message arrives. Such a 'Looped Receive' helps to realize a collect scenario.

[0093] The parallel scenario covers the situation in which two receive steps are waiting for the same message of the same interface in parallel. The messages arrive at different points in times (t1 and t2) while both receive steps wait from the correlation activation time on. The semantics are as follows: the first message does not complete both receives steps; each receive step gets its own message, while the order can be arbitrary.

[0094] Having described receive messaging and steps, send messages and steps will now be described. The send step is used to send a message that is already available within the process. Executing a send step, the process
5 engine submits the (outbound) message to the pipeline for processing. The receivers for a message to be sent can be determined or specified via routing configuration in the directory, or directly by process definition in combination with the receiver determination step or using another
10 message with which to reply).

[0095] When the routing configuration is retrieved from the directory, the process name and the outbound interface are the sender information within the message. The message is submitted to the pipeline where the suitable routing
15 relations are evaluated to determine target interface(s) and receiver(s). Context objects may be used to distinguish from among different send steps on the same interface, as shown in FIG. 13. When receivers are specified directly by process definition, the receiver (i.e. business system) may
20 be entered directly or calculated by the receiver determination step. In this case, the pipeline must not calculate the receivers again upon the Routing configuration.

[0096] If the routing configuration (directory) is used,
25 context objects are needed to distinguish send steps that send messages of the same interface from different places in the process, e.g. the send steps are each located in an exclusive branch of a switch element and need to be sent to different receivers. Without context information, it may
30 not be possible to distinguish the send steps from each

other because they are part of the same process and are being sent on the same interface. Routing can only use the process name and interface name. Context objects will make it possible to submit additional context information via
5 send step to the routing. This information can be used in a routing condition and classify the routing relation true or false.

[0097] Asynchronous communication is the default behavior for a send step. The message is send to its receivers using
10 the pipeline (without receiving a response or a fault message). For asynchronous communication, it is possible to wait for acknowledgements (none/pipeline, technical, application (default: no acknowledgement)). These acknowledgements are no business messages but (positive or
15 negative) technical responses of the receiver. Acknowledgments can only be received, if the receiver (adapter, system, etc.) is able to send these acknowledgements.

[0098] Errors while sending will not be visible on the
20 process level. Accordingly, the send step does not return any error states. It ends when the message can be submitted to the pipeline and the process engine receives a requested acknowledgement (if any is requested). Retries are realized by the XI runtime (number of retries could be defined in
25 CPAs), as the process engine is not configured to execute any retries.

[0099] A send step can also be used for synchronous communication. In this case, the send step deals with two messages: one message that is sent and another message that
30 is received as a result. Using synchronous communication,

it is not necessary to handle acknowledgements since a call will immediately return a result.

[00100] The process engine is also able to send positive or negative acknowledgements via send step.

- 5 Acknowledgements that are requested by the caller systems and have not been answered explicitly will be satisfied when the process terminates normally.

- [00101] The receiver determination step calls the receiver determination of the pipeline to calculate the receivers of
10 a message using the routing information. The step returns a list of receivers. The list of receivers allows a process-controlled multicast. Like the send step, the receiver determination step provides context information that can be used for evaluating routing conditions.

- 15 [00102] A transformation step is configured to: merge messages (N:1); split messages (1:N); and transform one message from one interface to another (1:1 as special case). For each transformation step, it may be necessary to specify both the mapping program (interface mapping) and the
20 messages that take part in the transformation scenario. The mapping programs need to be able to handle the different transformation cases. The transformation of the messages is done by a service that is called by the process engine. The transformation service should be able to process complete
25 messages (payload, attachments, etc.).

[00103] Attachments that are linked to the message payload (e.g. catalog reference pictures as attachments) need to be treated by this service as follows. In a merge case, as illustrated by FIG. 14, all attachments are collected and

attached to the merged message. Attachments carry a name which is referenced by the message payload. All attachments must use a unique name (e.g. GUID). In the split case, all attachments can only simply be replicated and attached to
5 the created messages, even if a large amount of data is transported. Thus, the transformation step includes the operations of transform, merge and split. An operation of creating new messages is special case of the transformation step.

10 [00104] In operation, the transformation step calls a mapping service of the XI runtime, with one or several messages (not payloads) and the mapping program as import parameters, and receives one or more new messages (not only payload) as an export parameter. The transformation step
15 operates on messages, and as such, other value sources or targets (e.g. simple XSD types) are preferably not used.

[00105] Value mapping will not be possible within the process, as sender and/or receiver information that is needed for value mapping is not available within the process
20 and cannot be submitted to the mapping program. Therefore a process needs a 'normalized' message format regarding values. As shown in FIG. 15, the value mapping to the process chooses the values of sender business system A (BSA) as a process-normalized value-format. However, when a
25 process receives or sends a message, value mapping is possible.

[00106] Control flow relevant modeling elements or steps include switch, assign, control, block, parallel sections, loops, wait, empty, exceptions, and deadlines. A switch
30 provides several branches, each equipped with a condition

and an explicit order for evaluation. The first branch whose condition returns true will be taken at runtime. All other branches can be ignored. If no branch matches (i.e. all branch conditions return false), a default branch will be taken. The assign step changes the values of container elements (i.e. top level access). A message (payload, header, etc.) cannot be changed by this operation. To change the payload, the transformation step must be executed. The message payload can be accessed via Xpath, context objects, or other similar means.

[00107] The control step offers a number of different functions that influence the control flow, e.g. throw exceptions or terminate the process. The control step is a step type that is specialized for the different functions. The control step can include a throw exception function and/or a cancel/terminate function. In fault situations, it should be possible to throw an exception that will be handled by an exception handler. The exception that is thrown can be identified by the unique exception name.

In fault situations, it must be possible to involve an user at some point. This implementation of the control step provides the possibility to send an alert to SAP Alert Management.

[00108] The cancel/terminate functionality allows terminating a process and all participating steps without fault-handling and compensation behavior. Explicit and implicit cancel functionalities can be distinguished. The explicit cancel functionality of the control step terminates the whole process upon a certain process state (received message(s) or deadlines). Processing this step, the process

stops and the process instance itself as well as all active steps will be terminated..

[00109] The implicit cancel functionality is not modeled within a process, but is a feature of the process engine, as shown in FIG. 17. In comparison to the explicit termination, the implicit case does not affect the whole process but rather only a certain part of the process (e.g. one or more fork branches). One example of the implicit cancel functionality is a parallel section with two branches, where the first branch completes the parallel section at the join step (one of two). Accordingly, if one branch reaches the join step and the join condition is true, the other branch needs to be cancelled implicitly. Assume at t1 both branches are active. At t2 one branch reached the join step while the other branch is still in processing. At t3 the other branch is cancelled (join condition has been true for the first branch) and the process continues after the Join step. As with the explicit cancellation, no compensation is provided. The branch will be terminated implicitly without further notice. Eventual open communications must be considered when defining join conditions.

[00110] Blocks (known as "scopes" in BPEL or "sequences" in BPML) can be provided to define a collection of steps to be performed sequentially which can share a container definition. This fulfills also a documentation purpose. Blocks can also be used to define deadlines that cover more than one step, and to define exceptions and their reach. Since deadlines and exceptions will be defined for blocks,

it is also possible to define deadline and exception handlers at a block.

[00111] Exceptions can be raised using a control step, as described above, or by the system. The modeling tool provides the possibility to define exception handlers. Exceptions allow the process logic to signal a business fault. Several different exception types include application exceptions explicitly thrown by the process logic, and system exceptions thrown by the process engine. Exceptions can be defined at the block level. The block represents the scope of its exceptions. Each block may have an arbitrary number of exceptions. Each exception carries a unique name and may not be redefined.

[00112] Exceptions can be raised via the control step (throw exception) identifying the unique exception name. Exception handlers can be defined for the existing exceptions within their scope. Throwing an exception can be handled as follows. First, an exception handler is found for the exception name. The inner block that surrounds the throw command is first searched; if the exception handler cannot be found in the surrounding block, go to next higher block, and so on as long as the exception handler cannot be found (exception propagating). Second, all active steps within the block that carries the exception handler are stopped. If the handler is found in a parent block, this block with all active steps will be stopped. Third, the exception handler is started and the steps within the exception handler are executed. Fourth, at the end of the exception handler, the block (in which the exception handler is defined) is completed as normal. If an exception cannot

be handled in one of the blocks, the process will go to an error state.

[00113] Deadlines are provided to monitor due dates of a process, and to react upon a missed deadline. The modeling
5 tool also provides the possibility to model deadline handlers. Deadlines will only be available for the modeling element block; other elements will not support deadlines. This allows monitoring a set of steps and avoids race conditions. Race conditions could occur if the deadline is
10 attached to receive or send steps only and the deadline is reached before the step is instantiated.

[00114] Deadlines include the definition of the deadline due date, the deadline scope (when the deadline is activated and/or deactivated), and the reaction upon a missed
15 deadline. The deadline definition can be programmed not to respect weekends, public holidays, etc. If a due date is reached, a deadline exception will be raised that terminates all active steps within the block and starts the corresponding deadline handler. A corresponding deadline
20 handler can model the reaction upon a reached deadline. There may be no default deadline reaction. Each block may only carry one deadline. This deadline must have its own assigned exception name that can be used as a deadline exception. The deadline handler is defined for this
25 exception.

[00115] Deadlines include the deadline definition (date and time) and a reaction upon a missed deadline. The system supports definition of time restrictions for the execution of process activities (e.g. process has to STOP three days
30 before a payload specified deadline). The reaction can

entail sending a message to the appropriate business system, and/or an alerting mechanism (user notification e.g. via SAP Alert Management) for exception situations, e.g. if conditions are not fulfilled within a given timeframe.

- 5 **[00116]** The process builder is configured to provide two kinds of parallel sections: a static and a dynamic variant. The static parallel section (fork and join) provides a parallel section with an arbitrary number of branches, which number can be defined at definition time, and
- 10 synchronization. This parallel section will terminate upon the following conditions: as a specialized condition, a "n-out-of-m" logic is provided (all logic is also possible using "n-out-of-n"); as an arbitrary condition operating upon messages and container variables. Both conditions will
- 15 be checked whenever a branch reaches the join step. If one condition returns true, all other active branches will be terminated. The process continues with the next step immediately after the end of the parallel section.

- [00117]** The dynamic parallel section (ParForEach) is a
- 20 block that provides a parallel section with one defined branch and synchronization. The number of parallel executions of this branch can be determined at runtime. The dynamic parallel section has a multi-line element (table) assigned. The defined branch (the block) is executed for
- 25 each line of this element in parallel. If the multi line element is empty, the dynamic parallel section will be completed immediately. Each branch has its own address space. The dynamic parallel section will terminate upon certain conditions: as a specialized default condition, all
- 30 branches must complete; as an arbitrary condition operating

upon messages and container variables. Both conditions will be checked whenever a branch reaches its end (end of the block). If one condition returns true, all other active branches will be terminated. The process continues with the next step immediately after the end of the parallel section.

[00118] The Process Builder can provide different kinds of loops. One type of loop is an "while" loop. While loops repeat over a set of steps until an arbitrary condition is fulfilled. Another type of loop is a "ForEach" loop. A ForEach is a block that has a multi-line element (table) assigned. It loops over the given multi-line element and executes the associated steps for each line of the element in a sequence (one after another). If the multi-line element is empty, the ForEach loop will not be executed (i.e. the process will not go to error-state). The process continues with the next step immediately behind the end of the block.

[00119] The wait step specifies a delay within the process flow until a certain deadline is reached or a delay for a certain period has passed. The empty element or step simply does nothing and must act as a placeholder for steps to be modeled.

[00120] The control flow modeling elements having thus been described, modeling patterns will now be discussed. These patterns are high-level building blocks used to construct processes implementing the requirements as mentioned above. Each pattern can be combined with one or more other patterns and also with atomic process engine functions such as deadlines, exceptions, etc. Modeling

patterns include: serialize messages/sequences;
transformations/merge/split; multicast; and collect.

[00121] For messages and/or sequences to be serialized, several receive steps (to get all the needed messages into
5 the process) and several send steps can be combined in a manner which will guarantee the order. It is important that the send steps make use of the technical acknowledgements to ensure that a message actually drops into the receiver system. Alternatively, another receive step must be
10 modeled to get the corresponding business responses. Thus, such a pattern can include the steps of: receive messages; and send received messages in a given order respecting technical dependencies of receivers. For the latter step, wait for acknowledgement of a message before sending the
15 next message; and address one receiver per message.

[00122] Transformations/merge/split modeling patterns include: (N:1) transform collected messages to one new message (e.g. transform several invoices to one combined invoice); or (1:N) transform one message into several other
20 messages (transform a combined invoice to invoice respecting the original POs). The multicast modeling patterns include the steps of: receive message or create it via transformation; get receivers; and send out messages to receivers. The last step can be executed without caring
25 about answers/acknowledgements (Fire and forget), or by receiving a number of answers/acknowledgments (inverse serialization).

[00123] For the execution of message-relevant steps, the process engine relies on services that are provided by the
30 XI runtime and not by the process engine. The XI runtime

provides the following services used by the process engine:
send; receiver determination; and transformation. A
modeling tool can be provided. The process engine can also
provide a technical monitoring from message view (XI
5 runtime) and from process view (process runtime). It should
be possible to navigate from process monitoring to message
monitoring and vice versa. In addition, integration in the
XI monitoring infrastructure (i.e. runtime workbench) can be
performed. The process engine component test should signal
10 whether the process engine is alive. Finally, the processes
can transfer alerts in the CCMS-infrastructure.

[00124] Monitoring can be split up into two main areas:
monitoring of the overall state of the process engine; and
monitoring of one single process instance. For the former
15 area, the system supports process monitoring during runtime,
such as, for example, a graphical display of the actual
state of all active process instances. The system also can
support a data volume and runtime performance information
for the process monitoring. For the latter area, the system
20 is configured to write process logs for each process
instance (the application needs to write its own application
messages to the protocol), and to display process logs along
with the possibility to select with application data (e.g.
order number).

25 [00125] FIG. 16 illustrates an example of the business
process engine runtime 600. In the example, an SRM
application 802 communicates with a supplier application 804
via the integration server 206. Although represented in
runtime only with respect to SRM application 802, each
30 application includes inbound and outbound interfaces, and

may include a proxy for connectivity to the integration server 206. The integration repository 202 includes an instance of a business process 232 representing semantically-linked communication between the SRM application 802 and the supplier application 804.

[00126] The business process instance 232 includes a process definition that defines the communication, including logical and physical interfaces, between the SRM and supplier applications. The process definition defines the messaging action and content that takes place between services of those applications, and can be as fine-grained as desired.

[00127] The integration server 206 includes the business process engine 274, which receives and executes the business process instance(s) from the integration repository (and via the integration directory, which is not shown). The integration server 206 also includes messaging functions and services 806, such as logical routing, mapping, and physical address resolution. These functions and services 806 also include a map a transformation service, for merging, splitting, creating, and/or mapping messages between sender and receiver applications. The integration server 206 also provides core messaging services for Web-based messaging between application services.

[00128] Although a few embodiments have been described in detail above, other modifications are possible. Other embodiments may be within the scope of the following claims.

5

Claims

- 10 1. A system for integrating two or more applications in an enterprise system, the system comprising:

an integration server connected between the two or more applications, the integration server including a process engine configured to execute one or more predefined business processes; and

15

a runtime engine, operating under direction of the business process engine, for executing one or more messaging services between the two or more applications based on at least one predefined business process.

20

2. A system in accordance with claim 1, further comprising a repository for storing the one or more predefined business processes.

25

3. A system in accordance with claim 2, wherein the repository includes design-time objects that define interactions between the two or more application systems.

30

4. A system in accordance with claim 2 or 3, further comprising a directory for storing a portion of the one or more predefined business processes based on a runtime configuration of the enterprise landscape.

35

5. A system in accordance with claim 4, wherein the runtime portion of the one or more predefined business processes are accessible by the process engine for execution by the runtime engine.

6. A system in accordance with anyone of the preceding claims, wherein the messaging services include logical routing of a message, a mapping of the message from one application to another, and a physical address resolution of the two or more applications, in accordance with the one or more business processes executed by the business process engine.
7. The system in accordance with anyone of the preceding claims, wherein each business process defines message-based interactions between the two or more applications.
8. A system for message-based exchange between two or more business applications in an enterprise system, comprising:
- a repository storing enterprise system design-time descriptions of the two or more business applications, the design-time descriptions including design-time business processes; and
 - a directory storing enterprise system configuration-specific descriptions of the two or more business applications, the configuration-specific descriptions including configuration-specific business processes to be executed between the two or more business applications.
9. A system in accordance with claim 8, wherein the repository further includes design-time interface descriptions of the two or more business applications, message context descriptions of messages to be exchanged between the two or more business applications, and interface mapping descriptions associating pairs of the design-time interface descriptions, and wherein the design-time interface descriptions, the message context descriptions, and the interface mapping descriptions are linked to at least one design-time business process.
10. A system in accordance with claim 8 or 9, wherein the repository further includes design-time business scenario descriptions that represent an abstraction of the design-time business processes.

11. A system in accordance with claim 9, wherein each design-time interface description includes a link to at least one interface type object and/or at least one message type object.

5 12. A system in accordance with claim 9, wherein each design-time interface mapping description includes a link to at least one message mapping object defining a mapping between pairs of interfaces based on a business process.

10 13. A system in accordance with anyone of claims 8 to 12, wherein the directory further includes configuration-specific business scenario descriptions, which comprise the configuration-specific business processes.

15 14. A system in accordance with claim 13, wherein each configuration specific business process includes a link to at least one routing relation object and/or at least one mapping relation object.

15. A method for integrating two or more application services in a enterprise application landscape, the method comprising:

20 associating the two or more applications within a business scenario;

defining one or more design-time business processes based on the business scenario; and

25 linking each of the one or more design-time business processes to interface descriptions of the two or more applications.

30 16. A method in accordance with claim 15, further comprising linking each of the one or more design-time business processes to at least one context object that describes a routing of a message between two of the applications based on message content.

17. A method in accordance with claim 15 or 16, further comprising linking each of the one or more design-time business processes to at least one interface

mapping description that describes a mapping between interfaces of two or more applications.

- 5 18. A method in accordance with anyone of claims 15 to 17, further comprising defining one or more configuration-specific business processes associated with a runtime configuration of the enterprise application landscape and based on the design-time business processes.
- 10 19. A method in accordance with claim 18 further comprising:

generating an instance of at least one configuration-specific business process in a runtime engine; and

15 executing the instance to communicate messages according to the at least one configuration-specific business process.
20. A method in accordance with claim 19, further comprising communicating a message between two of the two or more applications according to the configuration-specific business process.

1/17




FIG. 1

2/17




FIG. 2

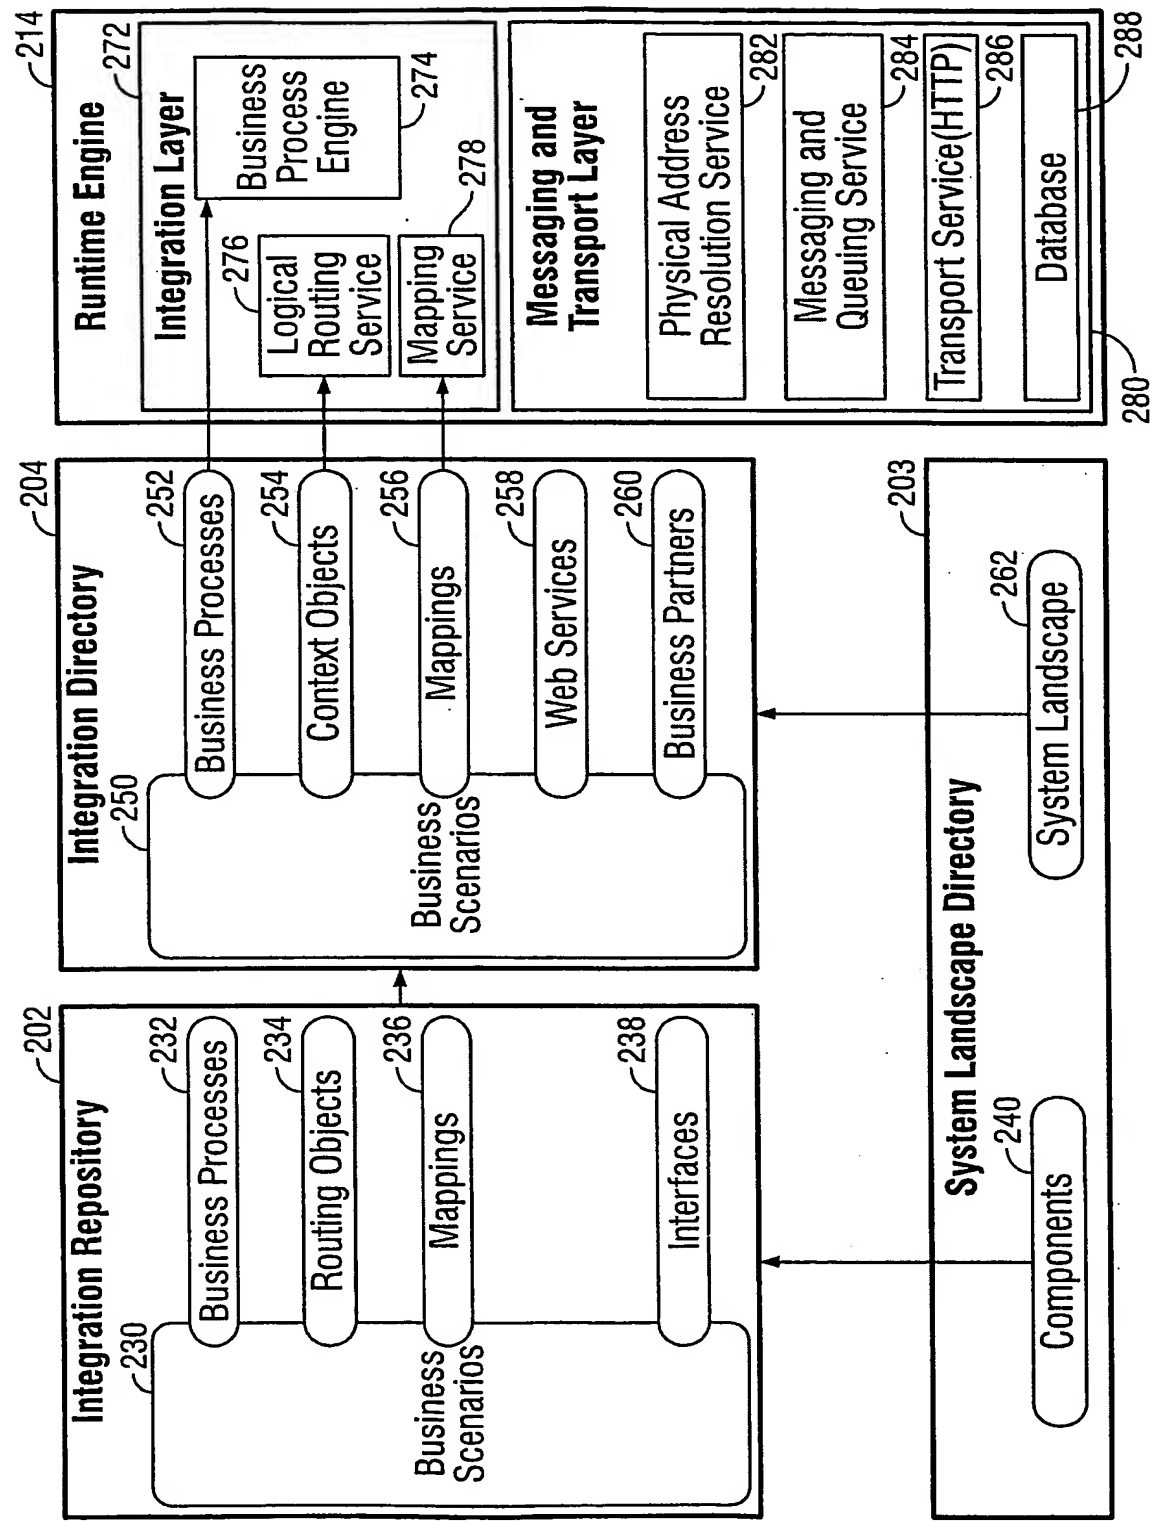


FIG. 3

4/17

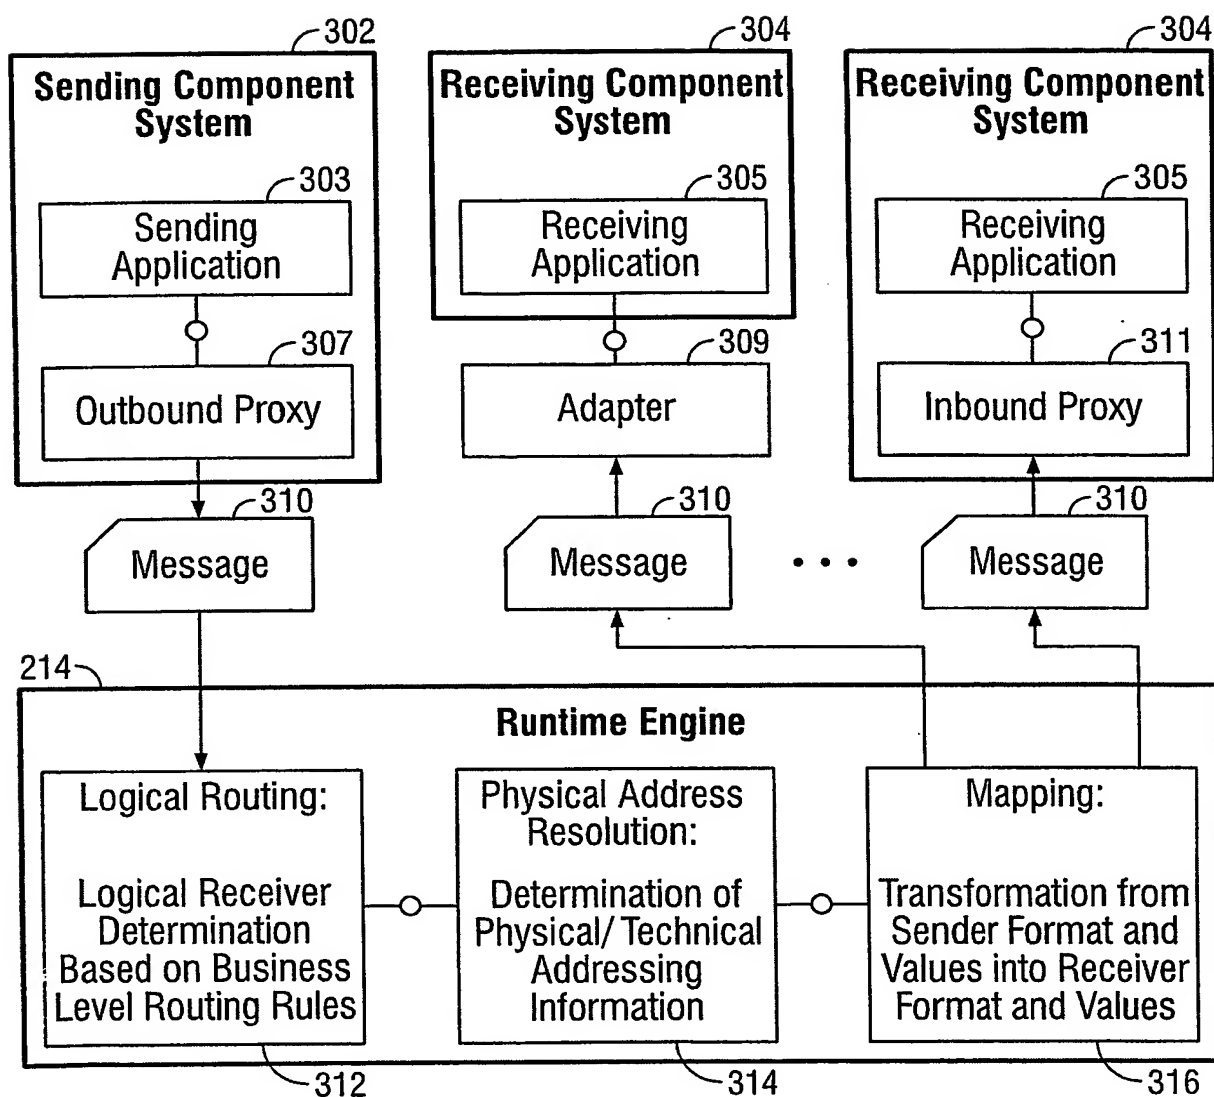


FIG. 4

5/17




FIG. 5




FIG. 6

7/17




FIG. 7




FIG. 8

9/17




FIG. 9

10/17




FIG. 10

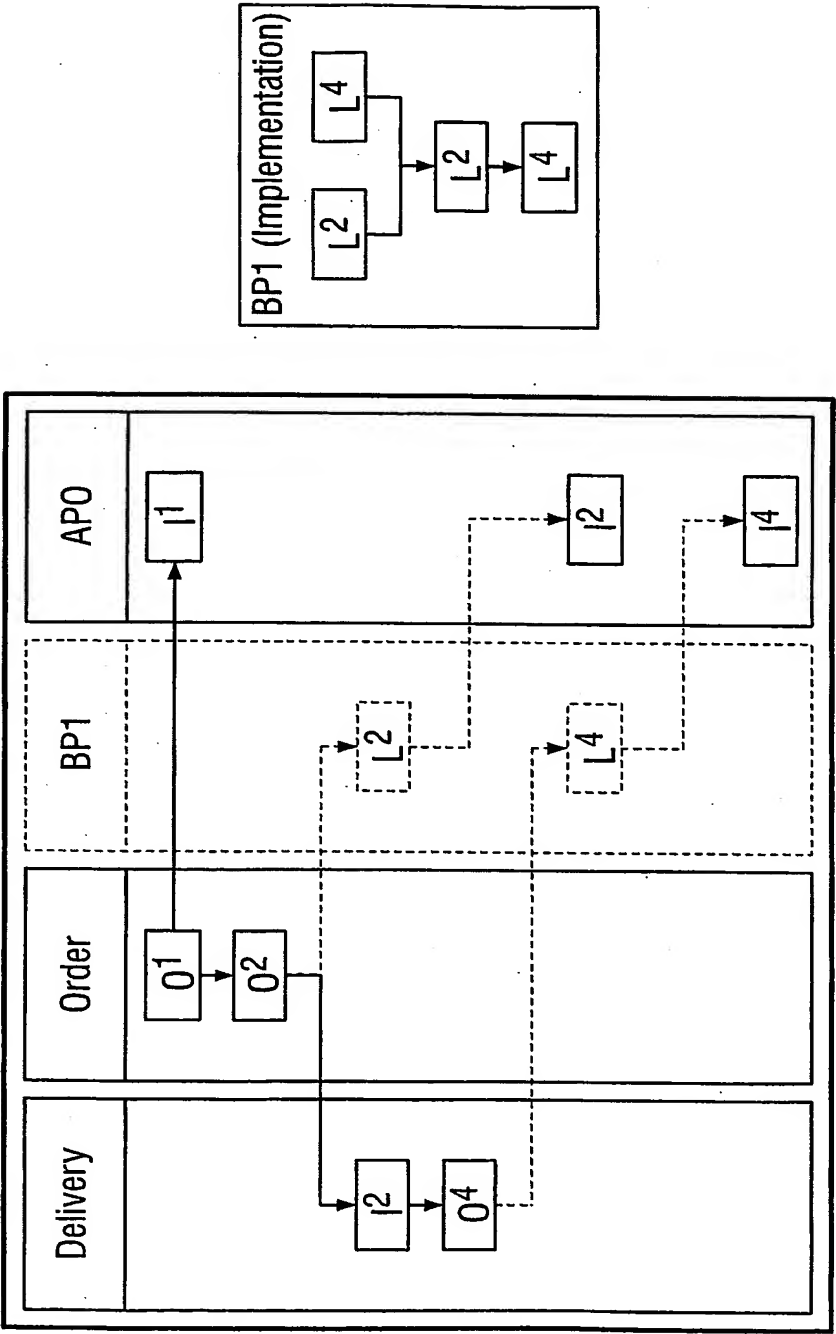


FIG. 11

12/17

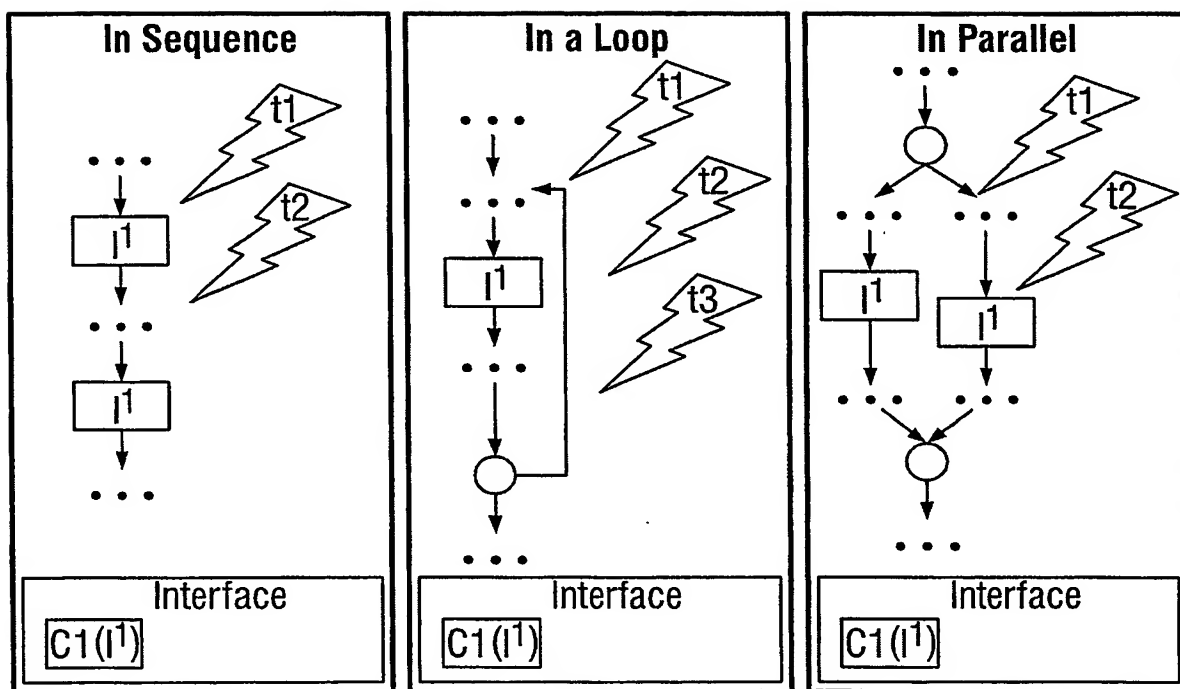


FIG. 12

13/17

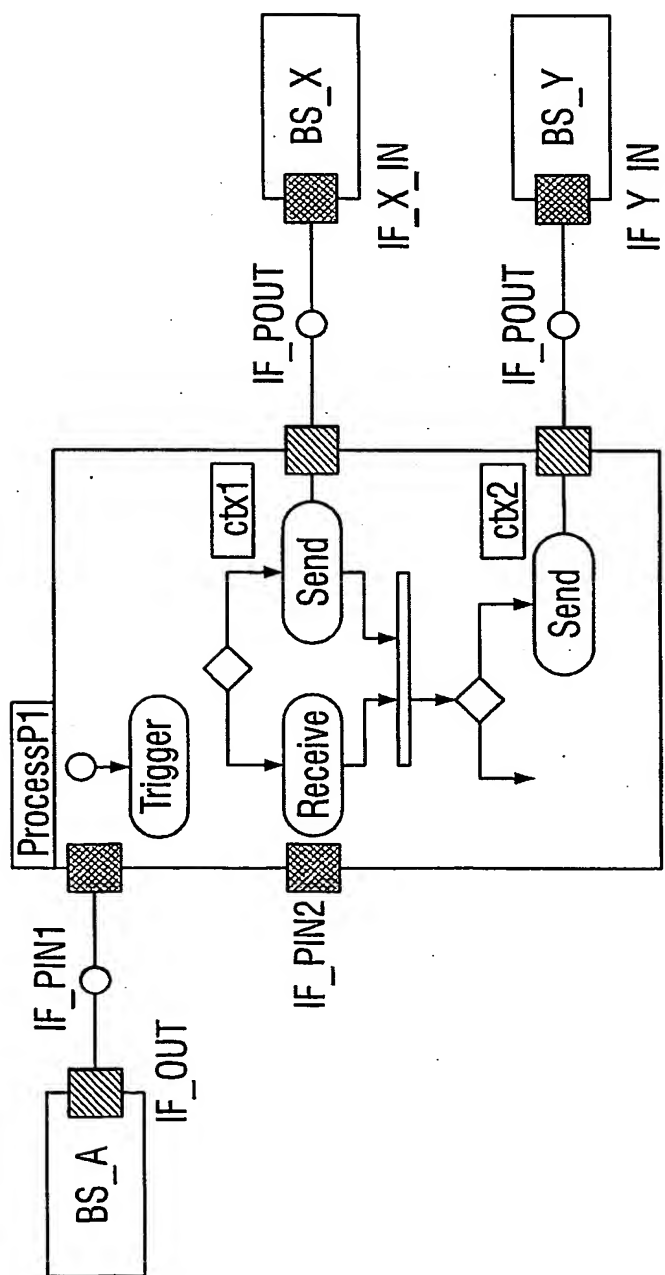


FIG. 13

14/17

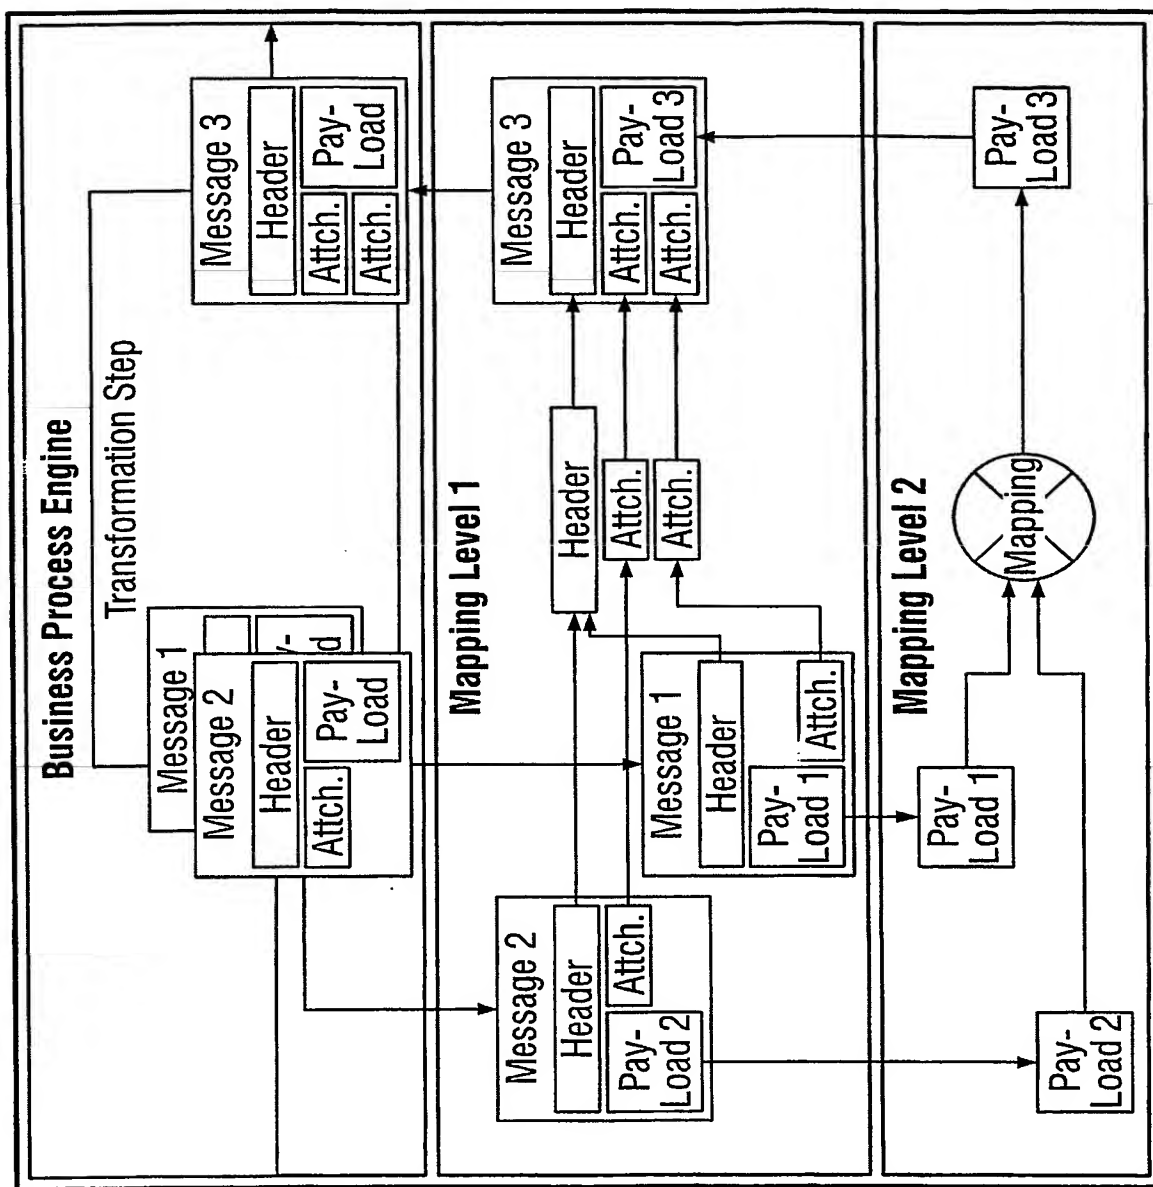


FIG. 14

15/17




FIG. 15




FIG. 16

17/17

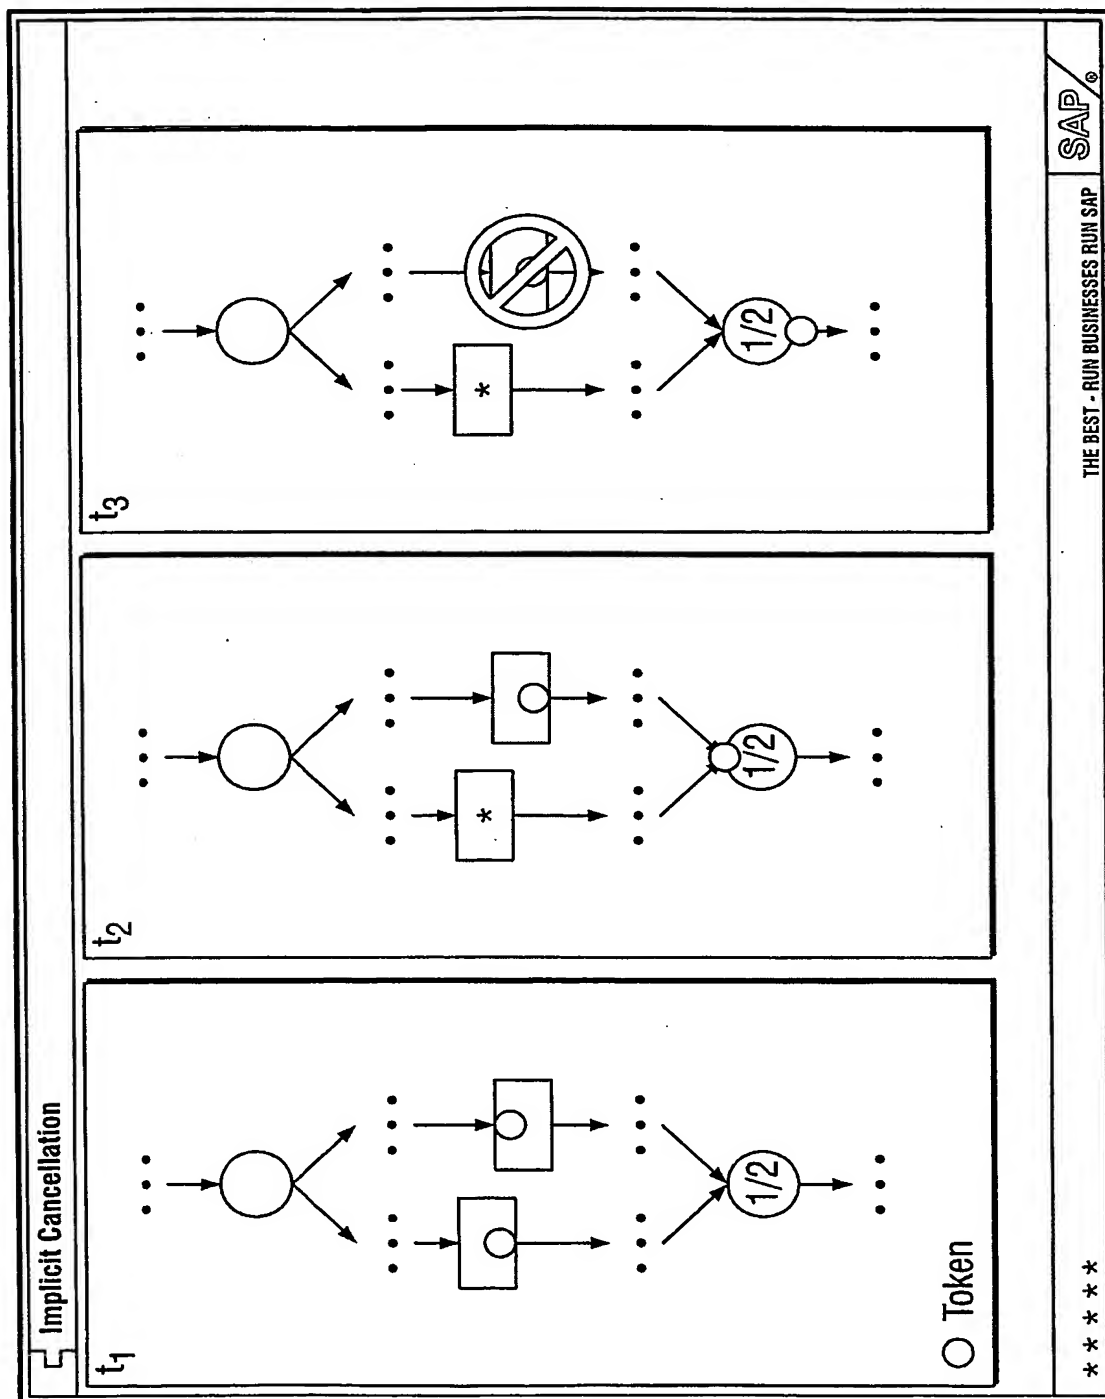


FIG. 17

THIS PAGE BLANK (USPTO)